The Deployment Gap

The AI selloff is pricing the financing, not the demand

On November 5 the semiconductor index fell 12.25%, its worst session since March 2020, and the S&P 500 had its worst day in eighteen months. The selloff came weeks after the strongest set of AI infrastructure results on record. Every layer of the stack that reported in the third quarter, from the model developers to the foundry, showed revenue growth and demand ahead of supply. The market's concern is real, but it is about who pays for the buildout and when it earns a return. It is not about whether the demand exists. The deployment gap, the distance between what AI can already do and how little of it has been put to work, is still wide, and closing it is what drives the buildout.

BY ADAM PATTI, CHIEF EXECUTIVE OFFICER, VISTASHARES | NOVEMBER 2026

KEY TAKEAWAYS

1. Demand has not changed, and it is not confined to the U.S. Every layer of the AI stack grew in the third quarter, and suppliers say demand still exceeds supply. Sovereign programs in the Gulf, Europe and Asia add buyers beyond hyperscalers, and the Senate just advanced a bipartisan permitting bill. Where one jurisdiction says no, capacity moves to one that says yes.

2. The concern is financing. Capital spending exceeds operating cash flow at three of the four largest hyperscalers, and model developers still lose money on compute. The question is who funds the buildout and when it pays off, not whether customers want it.

3. The price has reset; the earnings have not. AI infrastructure trades near a market multiple while consensus growth is several times the market's. The gap is widest in the contracted, supply-constrained layers, where earnings are already visible.

1. THE ADOPTION CURVE IS EARLY

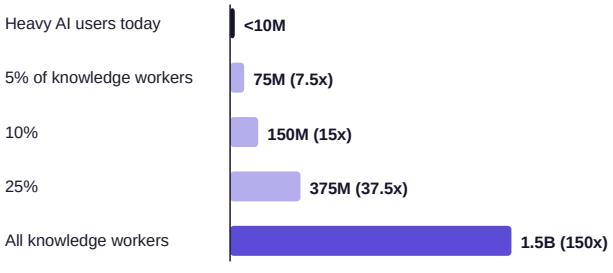
Capability is ahead of deployment. Venture firm Flybridge made this case to investors in October. Fewer than one in ten enterprises have AI agents in production, while current models can already handle much office and administrative work. Flybridge puts the remaining value from today's models closer to 100 times than 10 times. The constraint is deployment, not the technology. Usage is already compounding: Google's monthly token volume rose from 9.7 trillion in May 2024 to more than 3.2 quadrillion in May 2026, and API traffic rose another 40% last quarter.

The user gap is the market's blind spot. Gavin Baker of Atreides Management, whose figures Flybridge cites, estimates that fewer than 10 million heavy paying users support roughly $175 billion of annualized AI revenue, against 1.5 billion knowledge workers worldwide: less than 1% penetration. His own firm's token consumption rose 100-fold from March to August, and he says AI-native companies already spend more than 10% of payroll on tokens, against about 1% at the best traditional enterprises. In his words, the industry is "nowhere on the demand side" yet already supply constrained. A market pricing next year's financing is not pricing what happens as that penetration rises over the next five years.

The revenue data confirms the curve. Anthropic's amended S-1 shows third-quarter revenue of $16.96 billion, up 1,127% year over year and 47% from the second quarter. Over 11,900 customers spend more than $100,000 a year and generate 74% of trailing revenue, placing the growth in enterprise budgets. OpenAI told investors third-quarter revenue rose about 50% to roughly $10 billion, with enterprise run rate up 107%.

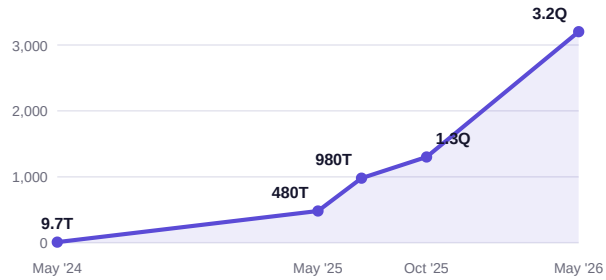
The user gap: heavy AI users

Scenarios vs. today, illustrative



Google: monthly tokens processed

Trillions, across all products; ~330x in two years



Left: heavy-user estimate from Gavin Baker, Atreides Management (a16z podcast, August 31, 2026); 1.5 billion global knowledge workers. Scenarios are illustrative, not forecasts, and revenue would not scale one-for-one with users. Right: Alphabet, CEO remarks at Google I/O (May 2024, 2025, 2026) and earnings calls (July and October 2025); T = trillion, Q = quadrillion. Token volume also rises as models reason at greater length, so it measures usage, not revenue.

2. EVERY LAYER OF THE STACK IS REPORTING GROWTH

LAYER	COMPANY	LATEST QUARTER	GROWTH	FORWARD SIGNAL
Models	Anthropic	$16.96B revenue	+1,127% y/y	$895B in compute and infrastructure commitments
Models	OpenAI	~$10B revenue	~+50% q/q	Enterprise run rate +107% in the quarter
Cloud	Google Cloud	$30.7B revenue	+103% y/y	~$561B backlog; $11.1B operating income
Cloud	Microsoft Azure	Not disclosed	+48% y/y	Capacity delayed to 1H27 by supplier constraints
Cloud	Amazon AWS	$46.0B revenue	+39% y/y	Fastest growth in nearly five years; $17.8B operating income
Compute	Nvidia	$96.2B revenue	+106% y/y	$108B guide vs. ~$104B consensus, ex-China; ~70% growth guided for FY28
Compute	AMD	$8.5B data center	+97% y/y	Q4 guided to ~$17B total revenue, ~+27% q/q
Foundry	TSMC	$46.7B revenue	+41% y/y	Q4 guided to $52.3B to $54.3B; demand still exceeds supply
Memory	Micron	$54B revenue (FQ4)	~5x y/y	75%+ of fiscal 2027 shipments under allocation or contract

Most recent reported quarter for each company. Nvidia figures are fiscal Q2 2027 (reported August 26); Nvidia reports fiscal Q3 on November 17. OpenAI figures are unaudited, as reported by The Information.

Spending is converting into revenue. The four largest hyperscalers guide to roughly $730 billion of 2026 capital spending, nearly double 2025, and consensus puts 2027 above $1 trillion. Unlike a speculative cycle, the cloud units selling this capacity are profitable and accelerating: AWS and Google Cloud earned a combined $28.9 billion of operating income in the quarter. Alphabet expects capex to "increase significantly" in 2027.

3. WHAT THE MARKET IS ACTUALLY WORRIED ABOUT

Funding, not demand. In the third quarter, capital spending consumed 68% of Microsoft's operating cash flow, 106% of Meta's, 121% of Alphabet's and 182% of Amazon's. Free cash flow at three of the four turned negative. Anthropic reported a $567 million operating loss after a profitable second quarter, as compute costs rose to 58% of revenue, and its $895 billion of commitments are roughly 13 times its current annual revenue pace. OpenAI's reported annualized revenue was revised from about $70 billion to about 50 billion once partner sales were counted on a net basis. These are legitimate questions about the financing structure of the buildout.

Rates and politics added pressure. October consumer prices rose 0.4% against a 0.3% forecast, and a December rate increase is now the base case. Higher rates raise both the discount rate on long-duration earnings and the cost of debt funding data centers. The midterm results added political risk around power prices. Neither changed the order books reported above.

4. THE BUILDOUT IS GLOBAL, AND IT GOES WHERE IT IS WELCOME

Washington is moving toward building, not blocking. Days after an election shaped by anger over power bills, the Senate voted 73 to 25 to advance the Bipartisan American Affordability and Jobs Act. The bill would speed approvals for power plants and transmission lines while requiring large data centers to pay more of the grid upgrades they require. Its Democratic co-author framed it as delivering on affordability and data center accountability. ING expects sweeping federal restrictions on data centers to remain unlikely while the current administration is in office. The political response so far is about who pays for the grid, not whether the capacity gets built.

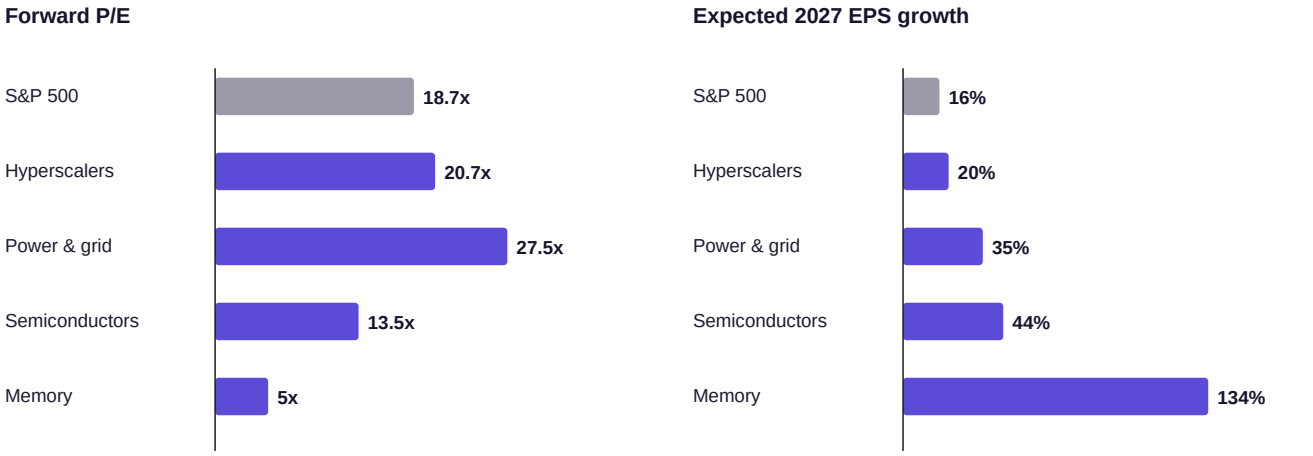
Hyperscaler capex is not the only buyer. Governments now treat compute as national infrastructure. Saudi Arabia's HUMAIN has committed roughly $100 billion across 11 data centers totaling 2.2 gigawatts. The UAE is building a 5-gigawatt AI campus with the United States, including the 1-gigawatt Stargate UAE cluster. The EU's InvestAI initiative targets €200 billion, including €20 billion for AI gigafactories, and Japan has pledged more than ¥10 trillion of public support for AI and semiconductors over seven years. Gulf commitments exceed $200 billion over five years. Some projects are still in procurement, but they add buyers whose demand does not depend on U.S. politics or hyperscaler cash flow.

Capacity relocates; it does not disappear. If a U.S. state restricts development, projects move to states competing for them or to countries with power, land and permits. For chip, memory, networking and power suppliers, where a data center is built matters far less than whether it is built.

5. WHY THIS IS THE ENTRY POINT

Supply, not demand, is the binding constraint. TSMC's chairman said AI demand remains above the company's ability to supply it. Microsoft pushed capacity into the first half of 2027 because suppliers could not deliver. Nvidia called its outlook supply-constrained and is raising prices to offset memory costs. Micron has most of next year's output already contracted. When customers are waiting for capacity, a debate about their long-run returns does not reduce near-term orders. It delays them at most.

Valuations have already reset. Before the November selloff, the median AI infrastructure stock had already compressed from 32 times forward earnings in April to 22 times, against roughly 19 times for the S&P 500, while consensus growth for the group runs from about 30% to more than 100%. The selloff took the semiconductor index roughly 30% below its high, and HSBC now puts U.S. semiconductors at 13.5 times forward earnings. Investors are paying close to a market multiple for several times the market's growth, and the discount widened while estimates held.



Forward P/E: S&P 500, FactSet; hyperscalers, average of Microsoft, Amazon, Alphabet, Meta and Oracle, S&P Global Market Intelligence (November 2); power and grid, average of GE Vernova and Vertiv; U.S. semiconductors, HSBC; memory, Micron on fiscal 2027 estimates. EPS growth: S&P 500, FactSet CY2027; hyperscalers, approximate, based on consensus estimates for Meta and group revenue growth; power and grid, analyst estimates of 32% to 39% for GE Vernova, Vertiv and Vistra; semiconductors, consensus as of mid-2026 via State Street; memory, Micron fiscal 2027 consensus. Bases differ by source; figures as of early November 2026.

The value capture sequence favors the suppliers. A capital cycle moves through income statements in order: the buyer's free cash flow compresses first, the supplier's revenue appears second, and the return on deployed capacity arrives last. The market is pricing stage one as a verdict on stage three. Companies that sell into the buildout, with contracted revenue and constrained supply, are where earnings are visible now.

6. POSITIONING AND WHAT WOULD CHANGE THE VIEW

Own the inputs, not the bets. The more durable exposure is the layers every model, cloud and enterprise deployment must pass through: leading-edge foundry, memory, networking, power generation and grid equipment. These layers carry multi-year contracts and pricing power, and they do not depend on which model developer wins. We would be more selective where growth depends on leverage or uncontracted capacity.

The signals to watch. Four data points will test this view over the next 90 days: Nvidia's results on November 17, the pricing and reception of Anthropic's public offering, the Federal Reserve's December decision, and hyperscaler 2027 capital spending guidance in late January and early February. A cut to 2027 capex plans, or a sustained drop in lab revenue growth, would weaken the case. Rising guidance with supply still constrained would confirm it. On today's evidence, the demand is real, the earnings are arriving, and the price of owning them has fallen.

SOURCES

Flybridge, "What We Told Our LPs About AI," Jesse Middleton (October 29, 2026). Gavin Baker, Atreides Management, on the a16z Podcast with David George (August 31, 2026), via BigGo Finance and Scuttlebutt. Google token volumes: Alphabet I/O keynotes and earnings remarks (2024 to 2026). Anthropic, PBC, Form S-1/A (November 2, 2026), via StockAnalysis, sec-api.io and FourWeekMBA. OpenAI third-quarter revenue: The Information via MLQ.ai and Crypto Briefing (November 10, 2026); annualized revenue revision: CNBC and Axios (October 8, 2026). Hyperscaler results and capital spending: company Q3 2026 releases and calls, CNBC (October 29, 2026), Walnut, ValueAdd VC, Ruslan Averin and Seat11a. Nvidia fiscal Q2 2027 results and guidance: company release (August 26, 2026), Zacks, SentiSense. AMD third-quarter 2026 results (November 3, 2026). TSMC third-quarter 2026 results, Form 6-K (October 15, 2026), Futurum, CRN Asia. Micron fiscal Q4 2026 results and earnings call (September 30 to October 1, 2026). AI infrastructure valuations: Goldman Sachs research via Bloomberg (September 25, 2026). Market data: Yahoo Finance and Reuters (November 5 to 10, 2026). October CPI: U.S. Bureau of Labor Statistics (November 10, 2026). U.S. permitting bill: Reuters, Washington Post and Associated Press (November 9 to 10, 2026). Sovereign AI programs: CNAS Sovereign AI Index, Presenc AI Sovereign AI Infrastructure Tracker, and Meticulous Research (2026). Capex-to-operating-cash-flow ratios are third-party calculations from company cash flow statements.

For informational purposes only. This material reflects the views of the author and is not investment, legal, or tax advice, nor a recommendation or offer to buy or sell any security. Companies are referenced for illustration and are not recommendations. Forward-looking statements and company guidance are subject to change. Past performance is not indicative of future results. All investing involves risk, including the possible loss of principal.